

What is it like to be a Dualist? On the real challenge of Nagel's 1974 paper

Lucas J Matthews

Philosophy

University of Cincinnati

1. Introduction

In his 1995 book, *The Engine of Reason, the Seat of the Soul*, contemporary philosopher Paul Churchland introduces his theory of recurrent neural networks. He argues that such networks provide a purely physical account of various mental phenomena. Being fully aware of the philosophical complexity of addressing the incessant and difficult “problem of consciousness,” Churchland seeks to defend the physicalist tradition from the famous anti-reductionist arguments offered by Thomas Nagel. In his classic 1974 paper, “What is it like to be a bat?” Nagel attacked physical-reductionism for its inherent inability to account for the subjective characteristic of consciousness.

In defense of his purely physical account of the brain and mind, Churchland attempts to account for the problematic subjective characteristic of consciousness. He seeks to reduce the subjectivity of experience to two purely physical modes of perception: sensory detection through *autoconnected* and *heteroconnected* neural pathways. I will address his account of unique neural pathways in this paper; however, it's important that I first introduce what follows next in Churchland's argument.

Although it would seem his reduction of the subjectivity of experience to purely physical neural pathways ought to properly deflect Nagel's attack, Churchland's vigorous defense of physicalism continues. He goes further to make the claim that Nagel's conclusions requires the existence of nonphysical objects: “. . . at some point . . . physical objects of knowledge are suddenly replaced by nonphysical ones . . . that is precisely what Nagel's conclusion requires” (1995, 199). Thus Churchland insults that Nagel's 1974 paper argues for dualism.

Churchland's accusation of dualism, however, is mistaken. Nagel's 1974 attack on physical reductionism does not provide an argument for dualism - a point that many readers seem to miss. In contrast, the often over-looked point of Nagel's argument against physical reductions focuses elsewhere: namely, to provoke the pursuit of a more objective understanding of the mental. More specifically, Nagel urged the development of a new objective phenomenology which might hope to capture the subjectivity of conscious experience. Thus, the often-missed point of Nagel's 1974 paper is develop a new method for filling-in where physical reductions run thin – *not* to argue for dualism.

If contemporary philosophers (such as Paul Churchland) *falsely* interpret Nagel's paper as an argument for dualism, this common misconception and the real point of Nagel's 1974 paper ought to be revealed. Thus, the main goal of this paper is to vindicate Nagel's 1974 from Churchland's accusation of dualism. However, questions still loom regarding what it is, precisely, that Nagel means, and hopes to do, with the development of "a new method and phenomenology." Therefore, in addition to a proper exposition and analysis of Nagel's anti-reductionist argument and Churchland's physicalist defense to it, my secondary goal in this paper is to investigate Nagel's idea of a new method and phenomenology and how it may be satisfied.

2. Background: Physical reductionism, dualism, and the problem of consciousness

For readers unfamiliar with the arguments at hand, I begin with a brief introduction to the fundamental ideas, concepts, and theories surrounding the disagreement between Churchland and Nagel. In the complex sphere of 'philosophy of mind,' two relevant explanatory traditions can be distinguished. The most common approach to questions regarding the philosophy of mind is to make an ontological distinction between what constitutes *the brain* and what constitutes *the mind*. This classical approach is founded in the theory of *dualism*, which essentially states that - in addition to the purely physical constituents of the brain - there exists something *nonphysical*. Traditionally, common examples of what might constitute the nonphysical entities are the mind, soul, spirit, energy, etc.

With the advent of new technology and increasing scientific understanding of nature, organisms, and the human brain, the dualist tradition is slowly diminishing. In its place, scientists and thinkers suggest that all things, including the mind, are of one single substance, and that substance is purely physical. This ontological approach to what exists is known as *physicalism* or *materialism*. What is likely to be the most common physicalist approach to explaining the world, including the mind and various mental phenomena, is known as *reductionism*. The agenda of the physical reductionist is 'to reduce' all things to their most basic physical constituents - which leaves no room for nonphysical ones. And it is precisely this, a physical reduction of the mind and various mental phenomena, which Paul Churchland hopes to achieve with his introduction of recurrent neural networks.

In *The Engine of Reason*, Churchland (1995) applies his novel theory of recurrent neural networks to the defense of physicalism. In particular, he describes how complex recurrent neural networks can account for many complex mental phenomena and functions such as taste, color, smell, and face coding for sensory representation, and more importantly for consciousness. As it is a result of technological and scientific advancement in the last 30 years, much of what Churchland develops is uncharted territory, and requires little philosophical persuasion. However, fully aware of its sheer complexity and unknown nature, Churchland approaches the particularly difficult problem of consciousness with a more philosophically rigorous style.

Even with groundbreaking technological advancements that give us insight into the brain, such as functional magnetic resonance imaging (fMRI), very little can yet be said about consciousness. Traditionally, the most difficult aspect of consciousness is its subjective character. There is something very unique about the subjective experience that seems to transcend purely physical explanations. One may endeavor to understand *how* we see by reading a chapter on the visual sensory system from a biology textbook. Understanding that we perceive our direct environment because our retinas detect light waves bouncing off objects in our foreground, and that the stimulation of photosensitive neurons in our retinas spark chains of synapses across neural pathways to our visual cortex, is one thing. However, such a chapter on the visual system does not explain *what it is like* to perceive visually.

It is on account of the subjectivity of visual perception that objective descriptions run thin. Not surprisingly, objective descriptions for any of the human senses are difficult. Describing, in purely physical terms, what it's like to smell and taste a freshly-baked blueberry pie would be like painting the rainbow with only two colors: an endeavor which would never do true justice to what a rainbow really looks like.

But the difficulties don't end here. If we find difficulties in describing the subjective experience of the sensory perceptions that most of us are capable of having (such as visual perception), then describing those of sensory experiences we are incapable of having (such as echolocation) ought to prove impossible. And this is precisely the point Nagel hoped to make by invoking the question: *What is it like to be a bat?* Although both echolocation and visual perception are methods in which organisms can detect physical objects in their immediate environment, the method and neural-wiring are completely analogous. We 'see' things by detecting light waves bouncing off objects in our foreground. Bats however, 'see' things by detecting *sound waves* bouncing off objects in their foreground.

Thus, Nagel's clever question brings home the point that there are some features of conscious experiences that seem utterly inexplicable in purely physical terms (albeit, *any* terms). Churchland exemplifies this point while introducing Nagel's argument: ". . . no matter how much one might know about the neuroanatomy of a bat's brain and the neurophysiology of its sensory activity, one would still not know *what it is like* to have the bat's sensory experiences" (195). In other words, no matter how much testing, poking, and prodding neuroscientists might do on a bat brain, they will never be able to describe, *what it is like* to be a bat. Nagel's question is meant to show that physical reductions of subjective experiences are exhausted by purely physical descriptions.

3. Heteroconnected and autoconnected neural pathways

Churchland's defense employs the concept of perception by means of heteroconnected and autoconnected neural pathways. Both forms of perception, one of which accounts for the subjective, occur in purely physical systems. Churchland explains that, as the brain develops, unique neural pathways are created as a causal response to various sensory inputs. Because every brain receives a unique set of sensory input, the causal neural pathways developed, at the level of synaptic connections and strengths, are completely unique. Churchland makes this point by explaining that ". . . each one of us, the bat included, enjoys a unique set of intimate *causal connections* to the sensory activity of one's own brain and nervous system" (Churchland 196). Therefore, we all have the ability to perceive our environment and bodies through neural pathways that are completely unique to us.

And in addition to having one-of-a-kind neural pathways, we also have means of perceiving things in our bodies through which no other brains have the ability to perceive. We have a *direct* neural connection to specific nerves in our bodies that no one else does. In explanation of this point, Churchland invokes an example that describes the experience of perceiving yourself blush. Imagine, for example, you're blushing because someone has complimented your pretty eyes. Both your admirer and you can perceive, by means of each of your own sensory systems, that you are blushing. Your admirer perceives it with her visual system; she can see your cheeks are red. You, however, perceive it by means of your somatosensory system; you can feel the embarrassing warm tingle reach across your face.

The sensation and perception of blushing on your face is purely physical: the microscopic capillaries in the epithelial tissue of your face are swelling with fresh oxygenated blood - which causes the tinge of red on your cheeks. The only difference in perception of the same objective phenomena (the occurrence of blushing) between you and your admirer is the means (neural pathway) by which you perceive the phenomena. You are sensing it through "autoconnected" means (neural pathways which are connected to perceive changes in your own body), while your admirer perceives it through "heteroconnected" means (neural pathways which are connected to perceive changes that are not occurring in your body). With this example, Churchland emphasizes that, although both perspectives may be different, the object of perception is the same - the physical occurrence swelling capillaries in your epithelial tissue: "the *object* of knowledge is exactly the same from both perspectives, the subjective and the objective, and it is something paradigmatically physical . . ." (197).

4. Why Nagel's 1974 paper is not an argument for dualism

Thus, it seems, Churchland has provided an account for how the unique subjective experience *can* be reduced to purely physical constituents: namely, perception of a phenomena through autoconnected neural pathways instead of heteroconnected neural pathways. However, as mentioned in the introduction, Churchland's physicalist defense

continues. He argues that, just because each one of us is able to perceive phenomena by means completely unique to us, and through which no others have access (autoconnected neural pathways), it does not follow that what is being perceived is nonphysical: “the existence of a proprietary, first-person epistemological access to some phenomenon does not mean that the accessed phenomenon is nonphysical in nature. It means only that someone possesses an information-carrying causal connection to that phenomenon, a connection that others lack” (1995, 198).

Churchland emphasizes that there is no need to posit nonphysical objects to aid in explanatory accounts of the subjective characteristic of consciousness: “There is no reason to expect, however, that this spectrum from knowledge to relative ignorance should reflect a hidden discontinuity at some point where physical objects of knowledge are suddenly replaced by nonphysical ones” (1995, 199). Churchland explains that, just because we run into difficulties describing the subjective phenomena, it does not follow that there must exist a mind or soul, or something nonphysical. But he insists, “this is precisely what Nagel’s conclusion requires” (1995, 199).

However, there is no mention of nonphysical objects or dualism in Nagel 1974. More importantly, Nagel doesn’t attempt to ‘fill explanatory gaps’ with nonphysical objects, nor does his conclusion require him to do so. Reasons for this reading are numerous. First, although Nagel argues for the existence of facts “. . . beyond the reach of human concepts,” nowhere does he make mention of nonphysical objects (1974, 4). Rather, he argues that it is perfectly rational to believe in the existence of ideas and facts that are beyond human comprehension and understanding: “Certainly it is possible for a human being to believe that there are facts which humans never *will* possess the requisite concepts to represent or comprehend” (1974, 4). But this doesn’t require that those facts include nonphysical substances or properties.

This fact ought not be difficult to concede. Consider the intellectual capabilities of a less-developed cognitive creature, such as a chimpanzee. It shouldn’t surprise anyone that such an animal lacks the capacity to comprehend organic chemistry. Actually, it shouldn’t surprise anyone that organic chemistry is something of which *most humans* lack sufficient understanding. Therefore, should it be difficult to accept that humans, also, have intellectual limitations? The possibility of the existence of things beyond the far reaches of our comprehensibility doesn’t seem completely unlikely. However, the truly pertinent point to draw is: even if we assume there exists things beyond our comprehensive capabilities, there is no reason to assume that the incomprehensible things are nonphysical by nature. Organic chemical substances and properties are not, obviously.

There are additional good reasons to believe that Nagel’s 1974 paper doesn’t intend to provide an argument for dualism. It does not follow that, if one denies the validity of one argument, then one must also accept the truth of the conclusion of its opposition. In Churchland and Nagel’s case, physical-reductionism and dualism seem to be the opposing theories. Nagel makes it very clear that he denies the explanatory capabilities of physical reduction, and thus its validity: “. . . we have at present no conception of what an explanation of the physical nature of a mental phenomenon would be” (Nagel 1). But does such a denial

imply an argument for dualism? Nagel makes no mention, nor argument, for dualism. Philosophers ought to approach all theories with skepticism, even the ones they support. This way, theories are constantly developed and their weaknesses are exposed for revision.

As an illustration of the above point, consider the agnostic who argues that atheism *cannot* be true. She might say: “*we cannot accept the conclusion of a valid argument as true unless we also know that its premises are true. However, the premise of atheism is that one knows God doesn’t exist. Proving the inexistence of God may be impossible, therefore we can’t accept the atheist’s conclusion as true.*” This is clearly an argument against atheism, but is it an argument for the opposing view, theism? The answer is *no*. The agnostic denies the truth of the conclusion of both atheism *and* theism. Is it not the case that one can deny physical reductionism *and* its opposition, dualism?

Finally, although Nagel’s 1974 ‘attack on physicalism’ is one key point of the paper, its larger spur is to provoke the development of ‘a new method and phenomenology.’ Nagel says explicitly that it would be a *mistake* to claim that physicalism is false (Nagel 7). What he really wanted to propose was that philosophers attempt to develop a new method and phenomenology for describing the difficult subjective character of subjective experiences. The goal would be “to describe, at least in part, the subjective character of experiences in a form comprehensible to beings incapable of having those experiences” (Nagel 9). This method and phenomenology would, for example, attempt to explain visual experiences in a manner that a blind person could comprehend.

What exactly does Nagel mean by “a new method and phenomenology?” As I hope to have shown, Churchland defended physicalism from an attack that Nagel never made; namely, that the inexplicable subjective character of conscious experiences provides an argument for dualism. However, if the real goal of Nagel’s 1974 paper was to provoke the development of ‘a new method and phenomenology for describing the subjective character of experiences in a form comprehensible to beings incapable of having those experiences,’ then a new question arises. Does Churchland’s appeal to recurrent neural networks take a step toward what Nagel 1974 was really after?

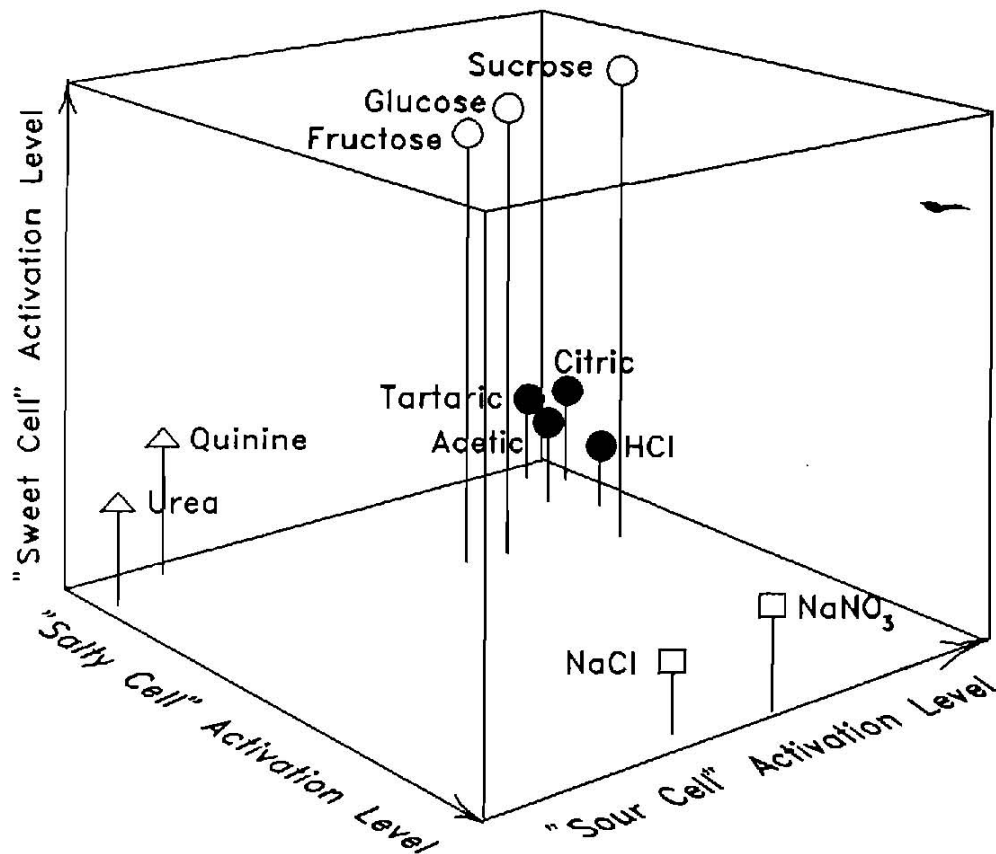
5. A possible new method and phenomenology: vector coding and three-dimensional “sensory spaces”

Surprisingly enough, although Churchland doesn’t explicitly attempt to provide a new method and phenomenology, he does provide a unique account for visualizing sensory representations. Perhaps his sensory state spaces are a step toward providing what Nagel is looking for. Churchland explains that the human capacity for verbal description using language is dwarfed by the incredible capacity of sensory functions because each relies on completely different coding strategies. In language, we (humans) use a finite domain of names to represent things, which often fall short of difficult descriptive tasks (like describing subjective experiences). In contrast, the human nervous system relies on an extremely

complex system of analysis and representation that Churchland explains in terms of *vector coding*.

Although Churchland's account of vector coding for various sensory experiences is far more complicated than I can explain here, the following is a brief account of its fundamentals. Churchland explains how one can evaluate the various synaptic weights of a network encoding a specific sensory experience – the taste sensation of something bitter for example – in a visually representational manner that takes the form of a three-dimensional diagram representing a specific 'sensory space.'

Figure 1 represents three of the four distinct types of taste receptors on the human tongue; sweet, sour, and salty (the fourth, bitter, is left out of the diagram for illustrative purposes).¹ When a substance is introduced to the taste receptors on the tongue, it consists of a specific combination of various chemicals that stimulate each distinct type of taste type differently. So an apple, for example, will heavily excite the sweet taste receptors, while the sour, bitter, and salty may receive a different combination of slight stimulation. (Churchland, for the sake of simplicity and illustration, assigns each taste receptor a level of stimulation on a scale from one to ten. Receptors at a stimulus level of one are activated much less by the substance than receptors at a level of nine or ten.)



Taste space: the position of some familiar tastes. (Adapted from Jean Bartoshuk.)

Churchland explains that each and every flavor, from sushi to bacon, has a distinct 4-element activation code or pattern – the combination of activities generated in each of the receptor types. The unique approach Churchland takes is to draw a three-dimensional diagram where each axis is represented by the activation level for each type of receptor. This creates what a sensory space, where specific flavors can be visually represented on the diagram. What's more interesting is that when this method is used, a natural organization of subjective taste experiences reveals itself. We find that particular corners and spaces of the sensory space cube constitute a specific type of receptor. This allows relative comparisons of various flavors pinpointed in the taste space.

All this become relevant toward considering how sensory space representations may be used to address Nagel's (1974) real challenge. Churchland covers four sensory modalities in his 1995 book: taste, color, smell, and facial recognition. It isn't difficult to see how these sensory spaces may be used to aid people incapable of having certain experiences. Although the color space may not help a colorblind person see color, it might certainly help them obtain an understanding of the relative relations of all the colors in the human visual spectrum. This is because each visual color experience corresponds to a specific activation pattern, in the retinal cone cells which can be mapped onto Churchland's color space. This way, the actual scientific configuration of real might be represented in a manner comprehensible to a colorblind person.

6. Final thoughts and conclusion

Such examples are few and still very conjectural. Philosophical applications of these ideas of sensory state spaces are still undeveloped, and a complete physical account of sensory inputs for each type of sensory experience will have to be established first. However, it's important to note that Churchland's theory *does* provide, although incomplete, a start toward describing experiences in a manner comprehensible to beings incapable of such experiences. However, the question still remains: how realistic is Nagel's 1974 request? Is a *full and true* new method and phenomenology possible? Is it necessary? What does Nagel hope to achieve with such an endeavor? Improving our descriptive accounts of subjective experiences seems a reasonable goal. However, filling in the content of his request in light of scientific understanding seems unclear.

One possible approach appeals to advances in technology. The fact of the matter is that science has been using technology to provide experiences humans have never been, and never will be, capable of having using their own natural sensory systems. A few moments on an internet browser generating images from specific technologies across the your computer screen: ultrasound, chemical-imaging, electron microscopy, fluoroscopy, magnetic resonance imaging (MRI), projection radiography, tomography, positron emission tomography (PET), photo acoustic imaging, etc. (the list goes on). Each of these technologies create visual representations of things we humans are incapable of experiencing unaided.

On reflection, however, this answer seems all-to-convenient. The real challenge lies in Nagel's original example: cross-sensory matching of analogous sensory systems. What's different about the above examples and describing what it's like to be a bat is that these machines conveniently translate things into a sensory function we are capable of in the first place. With echolocation however, we must attempt to translate experiences across sensory functions completely. But do we *need* to know what it's like to be a bat? Isn't, for example, ultrasound imaging close enough for Nagel's real purpose?

If we really wanted to challenge our scientific and philosophical capabilities, we might try translating specific sensory stimulus into sensory receptors not designed to receive such stimulus. For example, instead of translating all of these naturally invisible occurrences into visible images (like thermal radiation being translated into the human visual spectrum for example), one might endeavor to translate experiences of one existing sensory system, such as the gustatory system (taste), into a completely different system, such as the olfactory system (smell). Although completing such complex tasks may point us in the direct direction for answering Nagel's challenge, however, the question still looms what he's *really* asking for.

It would be necessary to complete a very scientific account of precisely what the human sensory systems are capable of detecting in the external environment. Briefly, the visual system can account for wavelength, frequency, and intensity of a segment of the electromagnetic radiation. The auditory system (hearing) detects changes of atmospheric pressure by registering wavelength, frequency and intensity of sound waves. The olfactory and gustatory systems give humans the ability to detect specific chemical signatures in the surrounding environment. Given the incredible speed in which we are advancing technologically, cross-translating a few of the above sensory inputs may not be so far out of sight.

And for Churchland's sensory space, such goals might still not be impossible. Consider a to-scale comparison of analogous sensory spaces. Where might the vector code of the color red fit on the taste space, or the smell space? Cross-translations and comparisons of Churchland's sensory spaces might be able to work as a bridge to a closer understanding of the relations of sensory experiences for people incapable of specific experience. Unfortunately for the blind however, Churchland's sensory spaces won't be so useful until someone devises a new method and phenomenology of creating sensory spaces representing at visual parameters that *don't* rely on a functioning visual system.

References

- Churchland, Paul (1995). The Engine of Reason; the Seat of the Soul. Cambridge, MA: MIT Press.
- Nagel, Thomas (1974). What is it like to be a bat? *The Philosophical Review* LXXXIII, 4: 435-50.
- Kawamura and Kare (1986). Umami: a basic taste: physiology, biochemistry, nutrition, food science. New York : M. Dekker.